

## **Supporting Information for**

Toward optimal disease surveillance with graph-based active learning.

Joseph L.-H. Tsui, Mengyan Zhang, Prathyush Sambaturu, Simon Busch-Moreno, Marc A. Suchard, Oliver G. Pybus, Seth Flaxman, Elizaveta Semenova, Moritz U. G. Kraemer

Corresponding authors: [lok.tsui@new.ox.ac.uk](mailto:lok.tsui@new.ox.ac.uk) (J.L.-H.T.), [elizaveta.p.semenova@gmail.com](mailto:elizaveta.p.semenova@gmail.com) (E.S.), [moritz.kraemer@biology.ox.ac.uk](mailto:moritz.kraemer@biology.ox.ac.uk) (M.U.G.K.)

### **This file includes:**

- Supporting Information Text
- Figs. S1 to S6
- Tables S1 to S6
- SI References

## Supporting Information Text

**S.1 Stochastic Susceptible-Infected Process on Graphs.** Each node is initially in the susceptible (S) state. At  $t = 0$ , a single node is randomly selected and set to the infected (I) state. At each following time step where  $t > 0$ , we assume that there is a fixed probability  $p = 0.1$  for any infected node to infect their susceptible neighbors, i.e. any nodes that are in the susceptible state and are connected to the infected node by an edge. In this simple SI process, it is assumed that infected nodes do not recover or become immune - once infected, they remain infected indefinitely with the same constant probability of onward transmission,  $p$ , throughout the rest of the simulation. This process continues until a certain proportion of nodes are infected, as specified by a target  $I/N$  value, where  $I$  is the number infected nodes and  $N$  is the total number of nodes in the graph. Different values of  $I/N$  indicate different stages of outbreak progression at the time of disease surveillance.

An important implication resulting from the assumptions made in the SI process as described, is that a node can only be infected if at least one of its immediate neighbors is also infected (with the exception of the initially infected node), i.e. all infected nodes must be connected in the graph. This implies that there can only be a single infected region, however with potentially multiple uninfected regions and therefore multiple decision-boundaries (lines or surfaces separating infected and uninfected regions). The distribution of these decision-boundaries in a network varies between outbreaks depending on both the network structure and the stage of outbreak progression (proportion of nodes that are infected).

**S.2 Conditional Autoregressive Model (CAR) as A Surrogate Model.** The Conditional Autoregressive (CAR) model (1) is widely used in the small area estimation domain, where data typically consist of observations  $\mathbf{y} = [y_1, y_2, \dots, y_n]$  over a set of  $n$  spatial units, which in the context of our study represent the observed infection status of a subset of locations in a mobility network. The CAR model assumes that the value of a variable at one location (e.g., infection status) depends on the values at neighboring locations, with weights specified by a spatial adjacency matrix  $\mathbf{A}$ . For unweighted models, such as the one we work with in this paper, the adjacency matrix  $\mathbf{A}$  is binary, capturing the presence or absence of edges between corresponding nodes. The spatial random effect  $\mathbf{f} = [f_1, f_2, \dots, f_n]$  follows a multivariate normal prior with precision matrix  $\mathbf{Q}$ :

$$\mathbf{f} \sim \mathcal{N}(0, \mathbf{Q}^{-1}) \quad [1]$$

$$\mathbf{Q} = \tau(\mathbf{D} - \alpha\mathbf{A}) \quad [2]$$

where  $\mathbf{D}$  is a diagonal matrix with diagonal elements corresponding to the number of neighbors each location has. The parameter  $\alpha$  captures the amount of spatial correlation between connected locations and can take any value between 0 and 1 (inclusive). If  $\alpha = 0$ , the model reduces to a set of independent errors at each location, whereas if  $\alpha = 1$ , the model reduces to the ICAR (intrinsic conditional autoregressive model) - another, but less flexible model. In this study, we fixed  $\alpha$  at 0.95 to clearly separate the task of spatial inference from the task of optimization, and use  $\tau \sim \logNormal(0, 0.1)$  as the prior on the marginal precision.

Much like Gaussian Processes (GPs) (2, 3) are a standard model of choice for continuous space modelling, CAR (and ICAR) are the default model choices for spatial statistics over a discrete areas. Future work should explore a wider range of surrogates, such as label propagation (4), Graph Neural Networks (GNNs) (5), and GPs on graphs when no knowledge about the spread of the disease is available, or spatially explicit mechanistic models like SIR and SEIR when the underlying transmission mechanics are known.

**S.3 Bayesian Active Learning by Disagreement (BALD).** Bayesian Active Learning by Disagreement (BALD) (6) is one of the state-of-the-art acquisition policies in active learning. It selects the data instances that maximise the decrease in expected posterior entropy,

$$v_{r+1} = \operatorname{argmax}_{v \in V} I(\theta; y|v, \mathbf{D}_r) \quad [3]$$

where  $\theta$  is the latent parameters and  $\mathbf{D}_r$  is the set of data instances labelled up to iteration  $r$ , and the mutual information,  $I$ , is defined as follows:

$$\begin{aligned} I(\theta; y|v, \mathbf{D}_r) &= H(\theta|\mathbf{D}_r) - \mathbb{E}_{y \sim p(y|v, \mathbf{D}_r)} H(\theta|y, v, \mathbf{D}_r) \\ &= H(y|v, \mathbf{D}_r) - \mathbb{E}_{\theta \sim p(\theta|\mathbf{D}_r)} H(y|\theta, v, \mathbf{D}_r) \\ &= \int p(y|v, \theta) p(\theta|\mathbf{D}_r) d\theta - \int H[p(y|v, \theta)] p(\theta|\mathbf{D}_r) d\theta \\ &\approx H \left[ \frac{1}{n} \sum_{i=1}^n p(y|v, \theta_i) \right] - \frac{1}{n} \sum_{i=1}^n H[p(y|v, \theta_i)] \end{aligned} \quad [4]$$

where  $\theta_i \sim p(\theta|\mathbf{D}_r)$ .

For Gaussian Process classification tasks, Houlby et al. (2011) (6) provided approximations of BALD. This formulation highlights that the mutual information can be approximated using posteriors obtained numerically. Hence, one can use surrogates of any complexity as long as their parameters can be estimated in a Bayesian manner.

**S.4 Selection by Least-Confidence (LC) and Node-Entropy (NE).** The policy Least-Confidence (LC) (4) selects the node for which the surrogate model is the least confident about its label prediction (label with the highest estimated probability among possible labels). More formally, this can be written as

$$v_{r+1} = \operatorname{argmax}_{v \in V} \max_{y \in Y} p(y|v, \mathbf{D}_r) \quad [5]$$

where  $Y$  is the set of possible labels,  $V$  is the set of unlabelled nodes available for selection, and  $p(y|v, \mathbf{D}_r)$  is the estimated probability of node  $v$  having label  $y$ , conditioned on the observed data up to iteration  $r$ ,  $\mathbf{D}_r$ . In the special case of binary classification where  $Y \in \{0,1\}$ , since the most likely label must by definition have an estimated probability greater than 0.5, this policy reduces to selecting the node with an estimated probability of having either label that is closest to 0.5. Without loss of generality, if we let  $p(y|v, \mathbf{D}_r)$  denote the probability of node  $v$  having label  $y = 1$  conditioned on the observed data up to iteration  $r$ , the policy LC can then be written as

$$v_{r+1} = \operatorname{argmin}_{v \in V} |p(y|v, \mathbf{D}_r) - 0.5| \quad [6]$$

The policy Node-Entropy (NE) (7) selects the node with the highest entropy in predicted label distribution, i.e.

$$v_{r+1} = \operatorname{argmax}_{v \in V} H[p(y|v, \mathbf{D}_r)] \quad [7]$$

where

$$H[p(y|v, \mathbf{D}_r)] = - \sum_{y \in \mathcal{Y}} p(y|v, \mathbf{D}_r) \log p(y|v, \mathbf{D}_r) \quad [8]$$

In the special case of binary classification, this expression reduces to

$$H[p(y|v, \mathbf{D}_r)] = H[p(v|\mathbf{D}_r)] = - p(v|\mathbf{D}_r) \log p(v|\mathbf{D}_r) - [1 - p(v|\mathbf{D}_r)] \log [1 - p(v|\mathbf{D}_r)] \quad [9]$$

with  $p(v|\mathbf{D}_r)$ , again, being the probability of node  $v$  having label  $y = 1$ , conditioned on the observed data up to iteration  $r$ .

From Eq. (9), it is straightforward to see that  $H[p(v|\mathbf{D}_r)]$  is a concave function of  $p(v|\mathbf{D}_r)$  that is symmetric about the line  $p(v|\mathbf{D}_r) = 0.5$ , i.e. when there is an equal probability of node  $v$  having either label. As a result, the node with an estimated probability  $p(v|\mathbf{D}_r)$  that is closest to 0.5 must also be the node with the highest entropy  $H[p(v|\mathbf{D}_r)]$ . Therefore, in the special case of binary classification, the policy LC always selects the same node as the policy NE at each iteration.

**S.5 Generating Random Graphs with Community Structure Using the Stochastic Block Model.** We used the stochastic block (SB) model (8) to generate random graphs with different levels of community structure. We began by first specifying the number of communities,  $k$ , and the size of each community. In this study, we set  $k$  to 5 with the size of each community selected at random while keeping the total number of nodes in the graph at 160. To control the level of community structure, we varied the value of parameters  $p_{intra}$  and  $p_{inter}$ , i.e. the probability of connection within a community and between communities, respectively. For example, a high  $p_{intra}$  with a low  $p_{inter}$  indicates a strong community structure, with nodes within communities being tightly connected and only sparse connections between communities. To generate a random graph with a high level of community structure, we set the parameters to  $(p_{intra}, p_{inter}) = (0.14, 0.001)$ ; and to generate a random graph with a lower level of community structure, we set  $(p_{intra}, p_{inter}) = (0.08, 0.005)$ .

One common way to quantify the level of community structure present in a graph is to compute its modularity (9). The modularity of a graph is a measure of the degree to which it can be partitioned into distinct modules or communities; it is defined as the fraction of the edges that fall within communities minus the expected fraction of edges that would fall within communities if edges were distributed randomly. Give a graph with adjacency matrix  $\mathbf{A}$ , its modularity is given by

$$Q = \frac{1}{2m} \sum_{ij} \left[ A_{ij} - \frac{k_i k_j}{2m} \right] \delta(c_i, c_j) \quad [10]$$

where  $m$  is the total number of edges in the graph,  $k_i$  and  $k_j$  are the degrees of node  $i$  and  $j$ , and  $\delta(c_i, c_j)$  is 1 if node  $i$  and  $j$  are in the same community, and 0 otherwise.

Applying the above formula shows that the random graph generated using the SB model with a high level of community structure has a modularity of 0.72, whereas the random graph with a lower level of community structure has a modularity of 0.62.

**S.6 Pre-Processing of Within-Country Human Mobility Data Collected at Provincial Level in Italy.** A dataset containing daily aggregated mobility data collected from mobile phone users in Italy at provincial level, covering the period between 18 January and 26 June 2020 (10), was downloaded from <https://covid19mm.github.io/data.html> on 26 February 2024. The data consists of the daily number of smartphone users moving both within and between 107 provinces, normalized by the number of active users each week, which has been shown to be roughly

constant throughout the collection period (3). Here we focus our analysis on the period between March and May (inclusive), during which a national lockdown (from 9 March to 18 May, 2020) was imposed by the Italian government in response to the emerging COVID-19 outbreak.

To construct a static graph from the mobility data, we first summed the mobility flows over both directions for each pair of provinces to obtain a symmetric matrix for each day, which was then averaged across the 3-month period. The resulting matrix was then converted into an unweighted graph using a procedure known as graph-thinning. In this process, edges representing pairs of provinces were ranked according to their total mobility flow as calculated earlier; edges were then removed one at a time starting from those ranked the lowest while ensuring that the graph remained connected. This iterative process continued until a certain target proportion of the original edges remained; this target proportion is known as the thinning-threshold. Finally, all edge weights are removed.

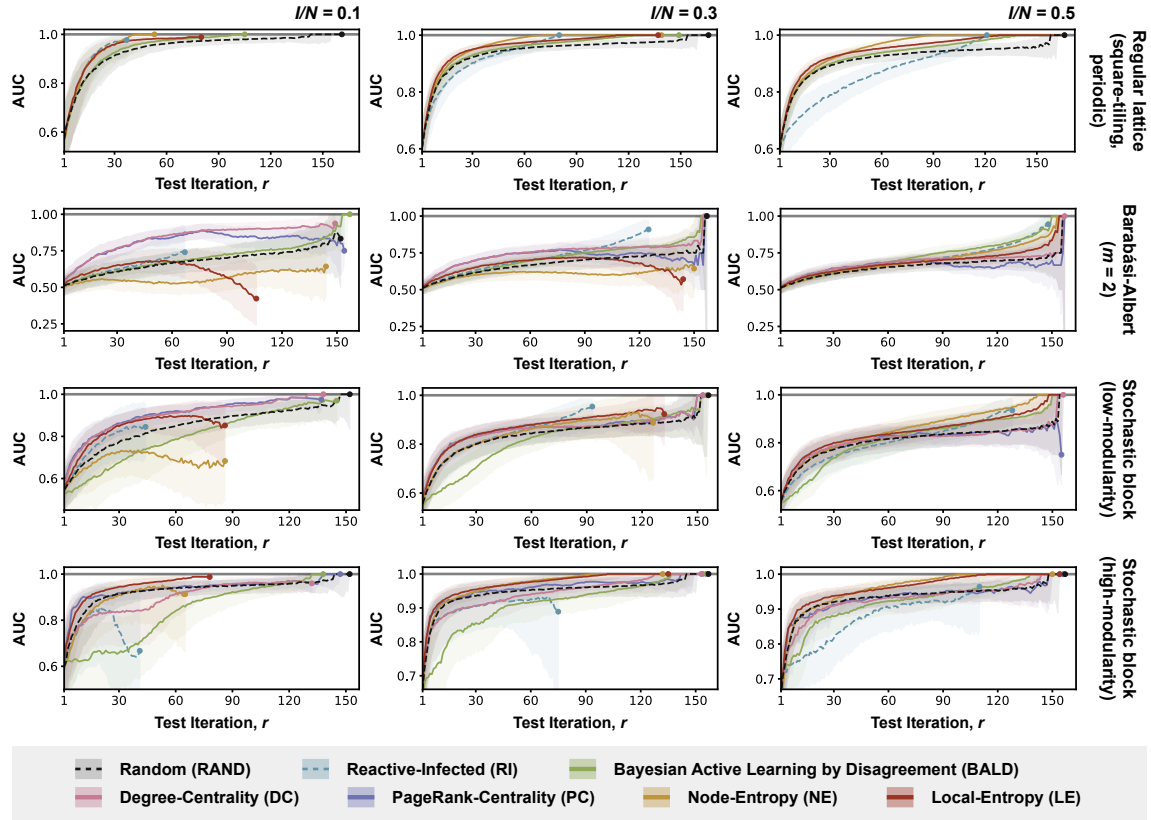
The choice of this threshold takes into consideration the balance between 1) the need to remove edges with very low mobility flows and are therefore less relevant to the overall structure of the graph, versus 2) the need to retain enough edges in order to preserve important structural properties (e.g., presence of travel hubs and community structure) of the graph. With these in mind, the thinning-thresholds of 10%, 15% and 20% were specified. To ensure robustness, the same experiments were repeated on each of the resulting graphs (see Fig. S3); however, only results from experiments performed on the graph with a thinning-threshold of 15% are presented in Fig. 4.

**S.7 Pre-Processing of Between-Country Air Traffic Data Collected at Country Level.** A dataset containing monthly air traffic data collected at country level, covering the period between January and March 2020 (11), was downloaded from <https://zenodo.org/records/7472836> on 1 March 2024. The data consists of the monthly number of air passengers travelling both within and between countries. To construct an undirected and unweighted graph from the data, the same procedure as described in S.6 was performed. Due to the much greater number of edges (as a result of a greater number of nodes and the presence of long-range movements in the air traffic network), a lower thinning-threshold was used to ensure the surrogate model can be fitted within a reasonable timeframe at each iteration given the available computational resources. With the considerations as described in S.6, the thinning-thresholds of 2.5%, 5% and 7.5% were specified. Again, the same experiments were repeated on each of the resulting graphs to ensure robustness of our results (see Fig. S4); however, only results from experiments performed on the graph with a thinning-threshold of 5% are presented in Fig. 4.

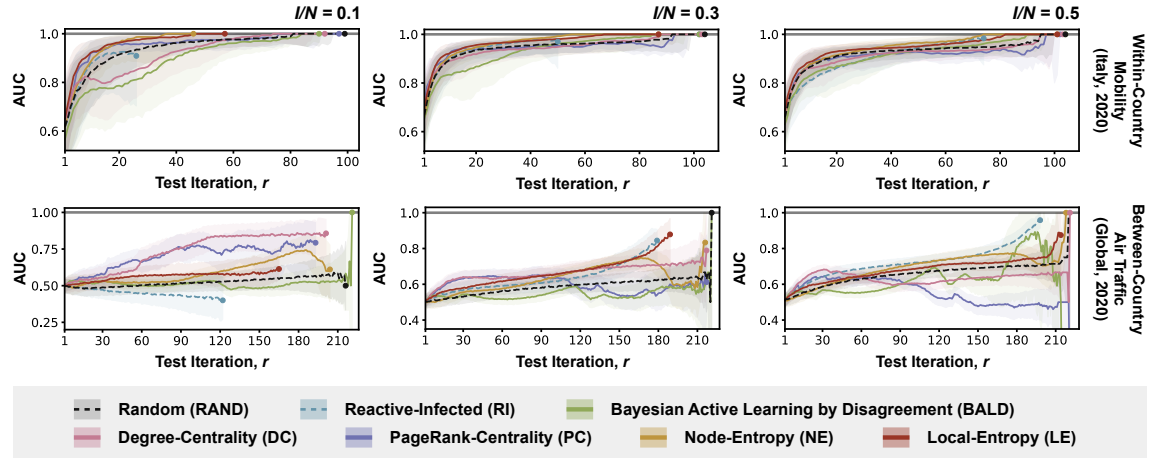
**S.8 Degree-Assortativity and Infection-Assortativity.** Degree-assortativity of a network, commonly denoted as  $r_{\text{degree}}$ , is a measure of the tendency for nodes to connect with other nodes with similar degrees. It can take any value between -1 and 1, with a positive value indicating that high-degree nodes are more likely to connect with other high-degree nodes, and similarly for low-degree nodes (assortative mixing by degree). Conversely, a negative value indicates a tendency for high-degree nodes to connect with low-degree nodes, and vice versa (disassortative mixing by degree).

The same idea of assortativity can be extended to other node attributes, including infection status as considered in this study. A positive assortativity by infection status (referred to as infection-assortativity hereafter) indicates a tendency for infected nodes to connect with other infected nodes, and similarly for uninfected nodes (assortative mixing by infection status). We denote the infection-assortativity of a graph with a given underlying disease distribution as  $r_{\text{infection}}$ .

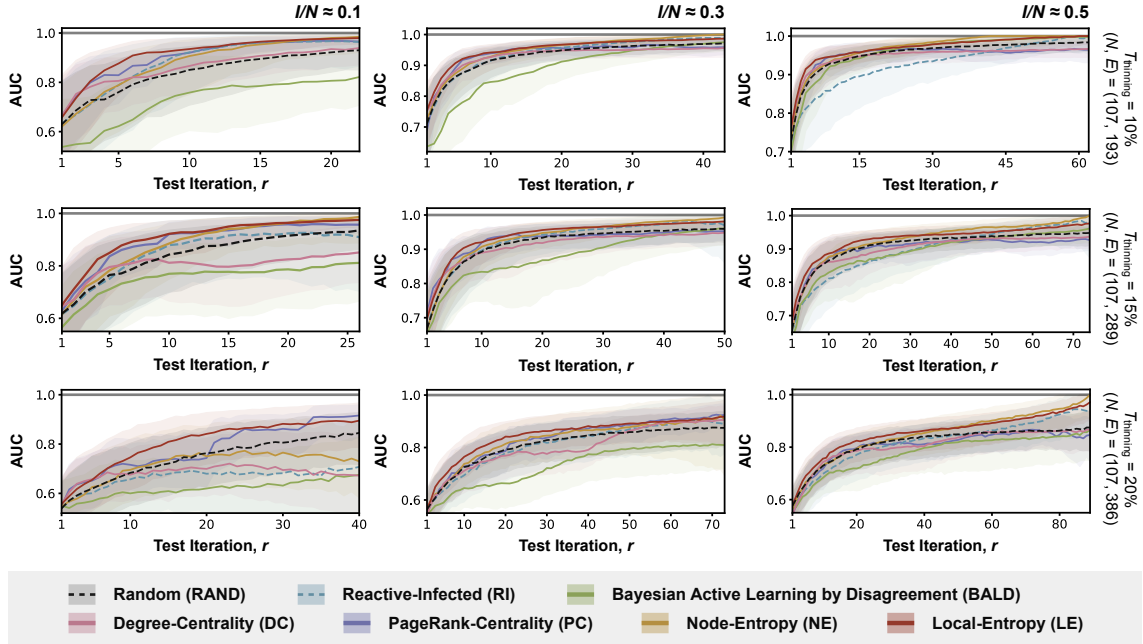
For a graph with an underlying disease distribution generated by a stochastic SI process (see S.1), we generally expect to observe a positive  $r_{\text{infection}}$ , since a node can only be infected if at least one of its immediate neighbors is also infected. The exact value of  $r_{\text{infection}}$  however depends on both the graph structure and the stage of outbreak progression (proportion of nodes infected) (see Fig. 3 and 4).



**Fig. S1.** Full results from experiments with simulated outbreaks on synthetic graphs. Each row presents results from experiments with simulated outbreaks on a different synthetic graph (as indicated by labels on the right); each column corresponds to simulated outbreaks at a different stage of outbreak progression, as measured by the proportion of nodes infected ( $I/N = 0.1, 0.3, 0.5$ ; as indicated by labels at the top). Each plot shows the performance of policies considered in the corresponding experiment, as measured by the AUC; the shaded region represents the interquartile range and the solid line represents the median. The performance of each policy is shown up to the median number of test iterations required for all infected nodes to be observed among agents assigned to that policy, with the AUC at this cut-off indicated by a colored dot (unlike Fig. 3, where the performance of each policy is only shown up to the median number of test iterations required for all infected nodes to be observed among agents assigned to Reactive-Infected (RI)).

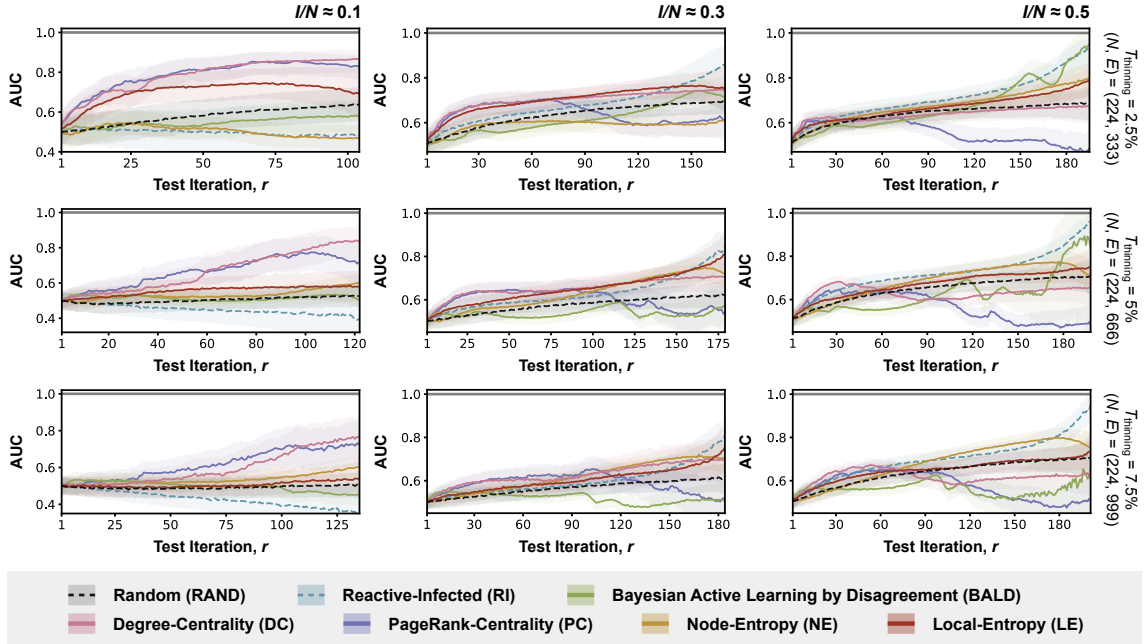


**Fig. S2.** Full results from experiments with simulated outbreaks on graphs derived from empirical human mobility data. Each row presents results from experiments with simulated outbreaks on a graph derived from a different empirical human mobility dataset (as indicated by labels on the right); each column corresponds to simulated outbreaks at a different stage of outbreak progression, as measured by the proportion of nodes infected ( $I/N = 0.1, 0.3, 0.5$ ; as indicated by labels at the top). Each plot shows the performance of policies considered in the corresponding experiment, as measured by the AUC; the shaded region represents the interquartile range and the solid line represents the median. The performance of each policy is shown up to the median number of test iterations required for all infected nodes to be observed among agents assigned to that policy, with the AUC at this cut-off indicated by a colored dot (unlike Fig. 3, where the performance of each policy is only shown up to the median number of test iterations required for all infected nodes to be observed among agents assigned to Reactive-Infected (RI)).

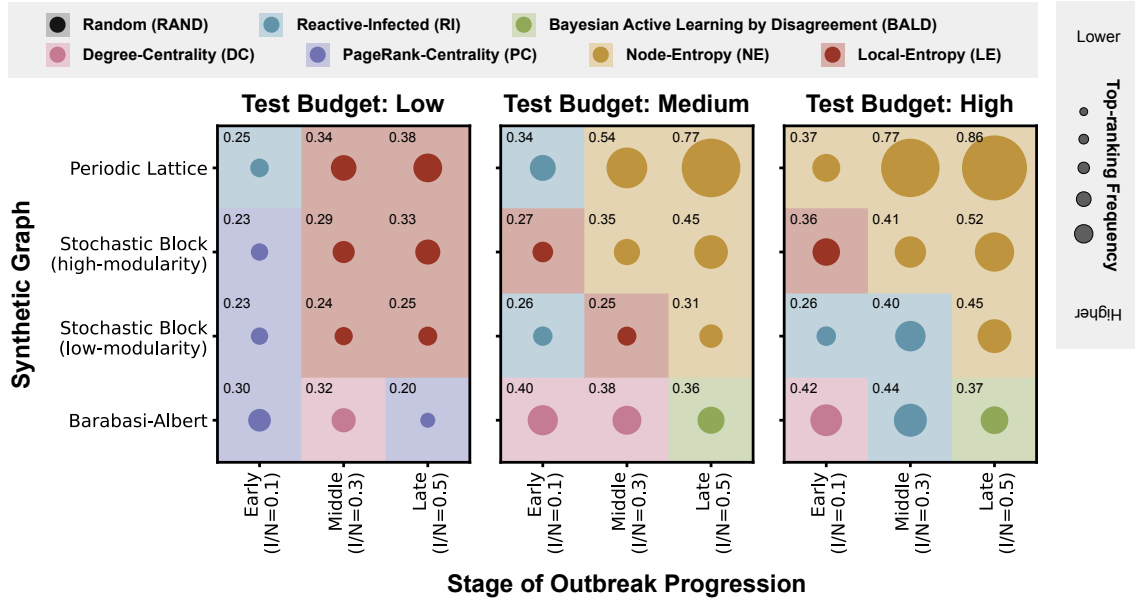


**Fig. S3.** Results from sensitivity analyses with simulated outbreaks on graphs derived from aggregated mobility data collected at provincial level in Italy. Each row corresponds to a different thinning-threshold ( $T_{\text{thinning}} = 10\%, 15\%, 20\%$ ; as indicated by labels on the right, with the number of nodes ( $N$ ) and edges ( $E$ ) remaining after graph-thinning also shown); each column corresponds to simulated outbreaks at a different stage of outbreak progression ( $I/N = 0.1, 0.3, 0.5$ ; as indicated by labels at the top). Each plot shows the performance of policies considered in the corresponding experiment, as measured by the AUC; shaded regions represent the interquartile range and the solid lines represent the median. Performance of each policy is only shown up to the median number of test iterations required for all infected nodes to be observed under the policy Reactive-Infected (RI).

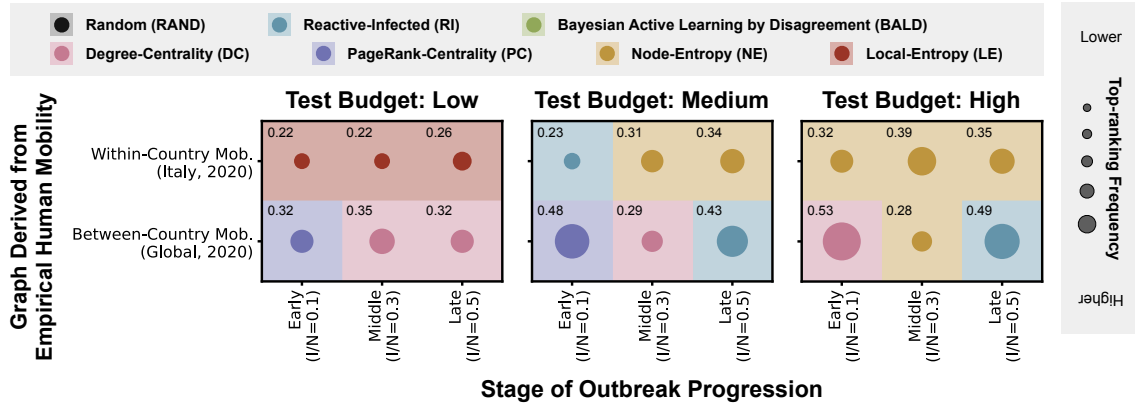




**Fig. S4.** Results from sensitivity analyses with simulated outbreaks on graphs derived from air traffic data collected at country level. Each row corresponds to a different thinning-threshold ( $T_{\text{thinning}} = 2.5\%, 5\%, 7.5\%$ ; as indicated by labels on the right, with the number of nodes ( $N$ ) and edges ( $E$ ) remaining after graph-thinning also shown); each column corresponds to simulated outbreaks at a different stage of outbreak progression ( $I/N = 0.1, 0.3, 0.5$ ; as indicated by labels at the top). Each plot shows the performance of policies considered in the corresponding experiment, as measured by the AUC; shaded regions represent the interquartile range and the solid lines represent the median. Performance of each policy is only shown up to the median number of test iterations required for all infected nodes to be observed under the policy Reactive-Infected (RI).



**Fig. S5.** Summary of top-ranking policies from experiments with simulated outbreaks on synthetic graphs. Each heat map presents results from experiments considering a different range of test budgets: low (up to one-third of the maximum test budget), medium (between one-third and two-thirds of the maximum), and high (above two-third of the maximum), with the maximum test budget determined by the median number of test iterations required by Reactive-Infected (RI) to identify all infected nodes (see Materials and Methods for more details). Each cell in a heat map corresponds to a specific synthetic graph (as indicated by labels on the left) and stage of outbreak progression (as indicated by labels at the bottom). The color of each cell indicates the policy that is most frequently ranked top across the test budget range, with the corresponding average top-ranking frequency represented by the size of the enclosed circle; the numerical value (to 2 decimal places) is shown in the top-left corner of each cell.



**Fig. S6.** Summary of top-ranking policies from experiments with simulated outbreaks on graphs derived from empirical human mobility data. Each heat map presents results from experiments considering a different range of test budgets: low (up to one-third of the maximum test budget), medium (between one-third and two-thirds of the maximum), and high (above two-third of the maximum), with the maximum test budget determined by the median number of test iterations required by Reactive-Infected (RI) to identify all infected nodes (see Materials and Methods for more details). Each cell in a heat map corresponds to a specific synthetic graph (as indicated by labels on the left) and stage of outbreak progression (as indicated by labels at the bottom). The color of each cell indicates the policy that is most frequently ranked top across the test budget range, with the corresponding average top-ranking frequency represented by the size of the enclosed circle; the numerical value (to 2 decimal places) is shown in the top-left corner of each cell.

**Table S1.** Top-ranking frequencies of policies in experiments with simulated outbreaks on a periodic lattice graph (with square-tiling) at different test budget levels.

* highest frequency: <b>BOLD</b> ** second highest frequency: <u>underlined</u>		Test Budget Level		
Stage of Outbreak Progression	Policy	Low (up to 1/3 of max.)	Medium (between 1/3 and 2/3 of max.)	High (above 2/3 of max.)
$I/N = 0.1$	LE	<u>0.219169</u>	0.2132	0.199100
	NE	0.203815	<u>0.2445</u>	<b>0.369300</b>
	BALD	0.164185	0.1051	0.072167
	RI	<b>0.246215</b>	<b>0.3431</b>	<u>0.302267</u>
	RAND	0.166615	0.0941	0.057167
$I/N = 0.3$	LE	<b>0.338489</b>	<u>0.273230</u>	0.063056
	NE	<u>0.248163</u>	<b>0.538637</b>	<b>0.771723</b>
	BALD	0.169348	0.085289	0.022918
	RI	0.093067	0.051911	<u>0.134338</u>
	RAND	0.150933	0.050933	0.007964
$I/N = 0.5$	LE	<b>0.380410</b>	<u>0.15483</u>	<u>0.061580</u>
	NE	<u>0.283737</u>	<b>0.77407</b>	<b>0.858502</b>
	BALD	0.163385	0.04440	0.017790
	RI	0.027424	0.00158	0.059115
	RAND	0.145044	0.02512	0.003013

**Table S2.** Top-ranking frequencies of policies in experiments with simulated outbreaks on a graph generated by the stochastic block model with high-modularity settings at different test budget levels.

* highest frequency: <b>BOLD</b> ** second highest frequency: <u>underlined</u>		Test Budget Level		
Stage of Outbreak Progression	Policy	Low (up to 1/3 of max.)	Medium (between 1/3 and 2/3 of max.)	High (above 2/3 of max.)
$I/N = 0.1$	LE	<u>0.203914</u>	<b>0.272657</b>	<b>0.363303</b>
	NE	0.123886	0.172429	<u>0.297600</u>
	BALD	0.040057	0.007029	0.020267
	PC	<b>0.229943</b>	0.151914	0.087590
	DC	0.106943	0.061343	0.034769
	RI	0.174371	<u>0.237314</u>	0.118328
	RAND	0.120886	0.097314	0.078144
$I/N = 0.3$	LE	<b>0.292496</b>	<u>0.318646</u>	<u>0.242381</u>
	NE	<u>0.188752</u>	<b>0.346246</b>	<b>0.414847</b>
	BALD	0.023328	0.002985	0.002750
	PC	0.183072	0.093969	0.055608
	DC	0.072512	0.035831	0.033936
	RI	0.103104	0.119338	0.216172
	RAND	0.136736	0.082985	0.034306
$I/N = 0.5$	LE	<b>0.330959</b>	<u>0.288663</u>	<u>0.264239</u>
	NE	<u>0.226094</u>	<b>0.445387</b>	<b>0.518526</b>
	BALD	0.030486	0.032541	0.044090
	PC	0.103805	0.037050	0.006027
	DC	0.105103	0.042423	0.015525
	RI	0.037773	0.079088	0.128737
	RAND	0.165780	0.074847	0.022856

**Table S3.** Top-ranking frequencies of policies in experiments with simulated outbreaks on a graph generated by the stochastic block model with low-modularity settings at different test budget levels.

* highest frequency: <b>BOLD</b> ** second highest frequency: <u>underlined</u>		Test Budget Level		
Stage of Outbreak Progression	Policy	Low (up to 1/3 of max.)	Medium (between 1/3 and 2/3 of max.)	High (above 2/3 of max.)
$I/N = 0.1$	LE	0.169067	0.192613	<u>0.209457</u>
	NE	0.100693	0.075760	0.087429
	BALD	0.040053	0.007093	0.008400
	PC	<b>0.228880</b>	<u>0.217120</u>	0.202543
	DC	<u>0.193867</u>	0.200907	0.195543
	RI	0.169707	<b>0.256640</b>	<b>0.257857</b>
	RAND	0.097733	0.049867	0.038771
$I/N = 0.3$	LE	<b>0.241110</b>	<b>0.252038</b>	0.162840
	NE	0.107935	0.150200	<u>0.225680</u>
	BALD	0.022271	0.009125	0.032627
	PC	<u>0.208942</u>	0.186113	0.073040
	DC	0.205652	0.144838	0.072720
	RI	0.125871	<u>0.210713</u>	<b>0.402493</b>
	RAND	0.088219	0.046975	0.030600
$I/N = 0.5$	LE	<b>0.251814</b>	<u>0.171265</u>	0.112317
	NE	0.154642	<b>0.310391</b>	<b>0.447413</b>
	BALD	0.039507	0.123619	0.137083
	PC	0.161470	0.111767	0.034381
	DC	<u>0.180130</u>	0.104000	0.021162
	RI	0.085674	0.114735	<u>0.222571</u>
	RAND	0.126763	0.064223	0.025073

**Table S4.** Top-ranking frequencies of policies in experiments with simulated outbreaks on a graph generated by the Barabási-Albert model at different test budget levels.

* highest frequency: <b>BOLD</b> ** second highest frequency: <u>underlined</u>		Test Budget Level		
Stage of Outbreak Progression	Policy	Low (up to 1/3 of max.)	Medium (between 1/3 and 2/3 of max.)	High (above 2/3 of max.)
$I/N = 0.1$	LE	0.119113	0.063418	0.041085
	NE	0.050487	0.007673	0.001855
	BALD	0.063043	0.016218	0.007291
	PC	<b>0.299600</b>	<u>0.386018</u>	<u>0.354448</u>
	DC	<u>0.296748</u>	<b>0.398673</b>	<b>0.419885</b>
	RI	0.106557	0.108782	0.165873
	RAND	0.064452	0.019218	0.009564
$I/N = 0.3$	LE	0.114038	0.075800	0.066634
	NE	0.049943	0.010457	0.012517
	BALD	0.045381	0.040610	0.151161
	PC	<u>0.315238</u>	<u>0.304905</u>	0.063405
	DC	<b>0.319629</b>	<b>0.382114</b>	<u>0.226098</u>
	RI	0.098819	0.159800	<b>0.435746</b>
	RAND	0.056952	0.026314	0.044439
$I/N = 0.5$	LE	0.150328	0.073935	0.055966
	NE	0.103792	0.105976	0.177720
	BALD	0.141528	<b>0.357869</b>	<b>0.366313</b>
	PC	<b>0.197864</b>	0.106792	0.031547
	DC	<u>0.190824</u>	0.148465	0.067289
	RI	0.116312	<u>0.152751</u>	<u>0.265550</u>
	RAND	0.099352	0.054212	0.035614

**Table S5.** Top-ranking frequencies of policies in experiments with simulated outbreaks on a graph derived from within-country human mobility data ( $T_{thinning} = 15\%$ ) at different test budget levels.

* highest frequency: <b>BOLD</b> ** second highest frequency: <u>underlined</u>		Test Budget Level		
Stage of Outbreak Progression	Policy	Low (up to 1/3 of max.)	Medium (between 1/3 and 2/3 of max.)	High (above 2/3 of max.)
$I/N = 0.1$	LE	<b>0.219754</b>	0.206427	<u>0.242877</u>
	NE	0.148102	0.219508	<b>0.318305</b>
	BALD	0.074824	0.022554	0.033323
	PC	0.158524	<u>0.225227</u>	0.166825
	DC	0.104304	0.034351	0.017657
	RI	<u>0.188401</u>	<b>0.229539</b>	0.164540
	RAND	0.106090	0.062394	0.056473
$I/N = 0.3$	LE	<b>0.219780</b>	<u>0.239300</u>	0.173316
	NE	0.159239	<b>0.311060</b>	<b>0.394824</b>
	BALD	0.048376	0.012331	0.045581
	PC	<u>0.162251</u>	0.060553	0.031986
	DC	0.154329	0.108904	0.045688
	RI	0.155200	0.203460	<u>0.270589</u>
	RAND	0.100824	0.064392	0.038016
$I/N = 0.5$	LE	<b>0.261104</b>	<u>0.217568</u>	0.170383
	NE	<u>0.196256</u>	<b>0.337395</b>	<b>0.348633</b>
	BALD	0.090016	0.118931	0.138441
	PC	0.132480	0.032816	0.011147
	DC	0.115040	0.069312	0.041220
	RI	0.066896	0.112088	<u>0.228010</u>
	RAND	0.138208	0.111891	0.062167



**Table S6.** Top-ranking frequencies of policies in experiments with simulated outbreaks on a graph derived from between-country air traffic data ( $T_{thinning} = 5\%$ ) at different test budget levels.

* highest frequency: <b>BOLD</b> ** second highest frequency: <u>underlined</u>		Test Budget Level		
Stage of Outbreak Progression	Policy	Low (up to 1/3 of max.)	Medium (between 1/3 and 2/3 of max.)	High (above 2/3 of max.)
$I/N = 0.1$	LE	0.131639	0.081161	0.048880
	NE	0.092790	0.026459	0.029920
	BALD	0.116537	0.012156	0.003260
	PC	<b>0.320390</b>	<b>0.478224</b>	<u>0.362887</u>
	DC	<u>0.220946</u>	<u>0.375717</u>	<b>0.526677</b>
	RI	0.061844	0.012937	0.023157
	RAND	0.055854	0.013346	0.005220
$I/N = 0.3$	LE	0.117387	<u>0.229433</u>	0.245539
	NE	0.033120	0.126787	<b>0.281207</b>
	BALD	0.033460	0.000773	0.003817
	PC	<u>0.330200</u>	0.194087	0.015254
	DC	<b>0.352613</b>	<b>0.294760</b>	0.193668
	RI	0.101347	0.132787	<u>0.249539</u>
	RAND	0.031873	0.021373	0.010976
$I/N = 0.5$	LE	0.104297	0.138991	0.072240
	NE	0.066545	<u>0.318454</u>	<u>0.242665</u>
	BALD	0.027958	0.026412	0.133206
	PC	0.135867	0.012967	0.000474
	DC	<b>0.321073</b>	0.009600	0.013846
	RI	<u>0.288491</u>	<b>0.431887</b>	<b>0.487163</b>
	RAND	0.055770	0.061690	0.050406

## SI References

1. J Besag, J York, A Mollié, Bayesian image restoration, with two applications in spatial statistics. *Annals Inst. Stat. Math.* **43**, 1–20 (1991).
2. Y. C. Zhi, Y. C. Ng, X. Dong, Gaussian Processes on Graphs Via Spectral Kernel Learning. Available at: <https://ieeexplore.ieee.org/document/10093993?denied> (Assessed 13 October 2024).
3. V. Borovitskiy *et al.*, "Matern gaussian processes on graphs". in *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR 130:2593–2601 (2020).
4. X Zhu, Z Ghahramani, "Learning from labeled and unlabeled data with label propagation" in *Technical Report CMU-CALD-02-107, Carnegie Mellon University* (2002).
5. T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks. arXiv [Preprint] (2017). <https://arxiv.org/abs/1609.02907> (Assessed 21 April 2024).
6. N. Houlsby, F. Huszár, Z. Ghahramani, M. Lengyel, Bayesian active learning for classification and preference learning. arXiv [Preprint] (2011). <https://arxiv.org/abs/1112.5745> (Assessed 12 December 2023).
7. D. D. Lewis, J. Catlett, "Heterogeneous uncertainty sampling for supervised learning" in *Machine Learning Proceedings 1994*, W. W. Cohen, H. Hirsh, Eds. (Morgan Kaufmann, San Francisco (CA), 1994), pp. 148–156.
8. PW Holland, KB Laskey, S Leinhardt, Stochastic blockmodels: First steps. *Soc. Networks* **5**, 109–137 (1983).
9. MEJ Newman, Modularity and community structure in networks. *Proc Natl Acad Sci U S A* **103**, 8577–8582 (2006).
10. E Pepe, *et al.*, COVID-19 outbreak response, a dataset to assess mobility changes in Italy following national lockdown. *Sci Data* **7**, 230 (2020).
11. S. Rudolf, Source for "How to avoid a local epidemic becoming a global pandemic?". *Zenodo* (2022). <https://doi.org/10.5281/zenodo.7472836> (Assessed 3 March 2024).